Proceedings of the First ISTA PDC Symposium on Seed Health Testing, Ottawa, Canada, August 1993

DESIGN AND INTERPRETATION OF SEED HEALTH ASSAYS

[1]Roberts, S.J., Phelps, K., Taylor, J.D., [2]Ridout, M.S.

[1]Horticulture Research International, Wellesbourne, Warwick, CV35 9EF, UK
[2]Horticulture Research International, East Malling, West Malling, Kent, ME19 6BJ, UK.

Summary

Disease avoidance by the use of healthy seed is a particularly important means of control for many bacterial and viral diseases. To monitor the production of disease-free seed, a wide variety of seed test assays have been developed and some are in routine use. The effectiveness of these assays is often limited by inappropriately designed testing schemes or incorrect interpretations of the results. We have attempted to combine practical considerations and recent mathematical studies to develop a rational approach to the design and interpretation of seed tests for bacterial diseases. The approach could equally apply to fungal or viral diseases in situations where it is necessary to detect low levels of seed infection. The merits of different test strategies are discussed in relation to tolerance levels and the establishment of confidence intervals, and in the context of routine commercial testing and epidemiological studies.

Introduction

Many important plant diseases are seed-borne. With the increased movement of seed around the world and the removal of trade barriers the need for accurate and reliable test methods has increased. This is particularly so when the pathogens involved are quarantine organisms. Bacterial and viral diseases present particular problems, as severe epidemics can result from relatively low numbers of infected seeds and, apart from the use of resistant varieties, there may be little prospect of control in the field. It is therefore essential that such diseases are controlled by a clean seed policy.

Numerous seed test assays have been developed for bacterial diseases. However, confusion frequently arises in the interpretation of such assays, especially in the case of negative results or when multiple samples of the same seed-lot are examined. Moreover, results are generally presented without any measure of precision.

This paper seeks to combine practical considerations and recent mathematical studies to develop a rational approach to the design and interpretation of seed tests for bacterial diseases both for routine testing of commercial seed and as a research tool in the study of disease epidemiology. The approach could equally apply to fungal or viral diseases in situations where it is necessary to detect low levels of seed infection.

There have been several previous papers on the design of seed health assays (Geng *et al.*, 1983; Russell, 1988; Marrou & Messiaen, 1967), but they have tended to concentrate on one type of assay design or one disease.

Test Methods

Methods for detecting seed-borne bacterial pathogens generally have a number of common features. Following sampling and, if appropriate, division into sub-samples, there is usually an extraction step, during which the pathogen is released from the seed into a liquid medium. The presence of the pathogen in the extract is then determined by immunological or molecular techniques or by the more traditional method of plating the extract onto selective or semi-selective agar media to produce bacterial colonies for further confirmatory identification. The

precise methods will vary between laboratories, and for different pathogens and seed types, and will be subject to different levels of sensitivity and detection thresholds.

Aims of the Assay

In designing a seed test, it is vital that the aims of the assay are clearly defined. Assays are of two basic types (qualitative and quantitative) which may be further sub-divided according to their purpose:

a) routine qualitative assays for quarantine purposes with a "zero tolerance";
b) routine qualitative assays for quality control with a pre-defined tolerance level;
c) routine assays for quality control where some quantitative estimate of the infection level is required (to make management decisions, risk assessments, etc.);
d) assays for research purposes, where precise estimates of infection levels are required.

The designs of assays for these different purposes are treated separately in the following sections.

Qualitative

The most important factor in the design of routine assays is the tolerance level, the maximum number of infected seeds which can be tolerated in a seedlot. There is no such thing as zero tolerance and it is impossible to certify that a seed lot has a zero level of disease. Unfortunately, tolerance levels have been determined for very few diseases, and it is often the case that tolerance levels arbitrarily arise as a consequence of the assay design. Setting of tolerance levels requires data on transmission rates, epidemiology and the relationship between disease levels and economic yield loss.

Seed tests for pathogens are usually expensive as they require a considerable amount of time, effort and materials to process each individual sample or sub-sample. Thus, an important feature of routine assays is that the number of samples to be tested needs to minimised. The design of the assay should be simple and interpretation should be straightforward. No assay is completely reliable and may give rise to two types of errors: either a seed lot with an infection level below the tolerance level is rejected (false positive) or a seedlot with an infection level greater than the tolerance level is accepted (false negative). The probabilities of these occurrences, usually given the symbols α and ß respectively, need to be considered in designing an assay. As will become apparent, it is also important to know the reliability with which a single infected seed can be detected in a sample of a given size or conversely the maximum sample size in which a single infected seed can be detected. There are also practical limits on the maximum sample size which can be processed in an assay.

The main differences between routine assays for quarantine purposes and quality control with a single pre-defined tolerance level will be the acceptability of false positive and negative results and the tolerance levels themselves.

Quantitative

The design requirements of quantitative assays differ from those above in that an attempt is being made to estimate the true level of infection in a seed lot. There may be no pre-defined tolerance level and the acceptability of false positives and negatives is no longer relevant. It is, however, important that there is some estimate of the accuracy and precision of the final

estimates obtained. There must still be some consideration of the reliability with which a single infected seed can be detected in a sample.

In routine testing the number of samples to be processed is still a consideration, but this may be less important for research purposes. Here, the time required for the final result may not be so critical and it may be feasible to do the assay in several stages over a period of several days or weeks.

Statistical Principles of Analysis

The theoretical basis for the analysis of seed test data is based on the key assumption that the seeds which are tested are a random sample from the seed lot. Each seed should be able to be classified as either healthy or infected and each seed (regardless of health status) should have an equal chance of being present in a sample. In practice, the mixing and sampling of seed lots to satisfy these assumptions can be problematical, but for the purposes of this paper it will be assumed to be adequate.

The probability, p_c of a least one infected seed being contained in a sample of size n from a seed lot containing a proportion of infected seeds θ is given by the binomial probability:

$$p_c = 1 - (1 - \theta)^n \qquad \textbf{(1)}$$

This model assumes that sampling is done with replacement, i.e. after the sample has been analysed it is replaced in the bulk. In practice this is not the case, but if the sample size is small in relation to the total size of the lot, it can be used as an approximation. This model is used for *limiting dilution assays* (Taswell, 1981), where the aim is to estimate the proportion of individuals in a population with some characteristic.

The Poisson probability model:

$$p_c = 1 - e^{-\lambda v} \qquad \textbf{(2)}$$

is used for *serial dilution assays* (Cochran, 1950) where the aim is to estimate the density of organisms in a liquid, where λ is the density and v is the volume sampled and is often known as the most probable number (MPN) method. This Poisson model has been used for seed health assays (Geng *et al.* 1983) but is not the same as the binomial model (eq.1). They can be made equivalent if we take $n = v$ and $1 - \theta = e^{-\lambda}$. It follows, therefore, that $\lambda = -\ln(1 - \theta)$ and when θ is small (< 10%) $-\ln(1 - \theta)$ is approximately equal to θ. Thus, the Poisson model is a valid approximation to the Binomial model when θ is small.

The apparent simplicity of the underlying principles and assumptions is misleading and the design of assays depends critically on their application as described hereafter.

Case 1: detection always possible, quarantine or quality control - pre-defined tolerance level

If it is always possible to detect an infected seed in a sample of a certain size, then the probability of a positive result, p_+, for a seed test on a sample will be the same as the probability of an infected seed being present in the sample, i.e. $p_+ = p_c$. It is therefore a simple matter to rearrange equation (1):

$$n = \frac{\ln(1 - p_c)}{\ln(1 - \theta_{nt})} \qquad \textbf{(3)}$$

to determine the size of sample which it is necessary to test in order to satisfy a set probability, p_+ (= p_c in this case), of a positive result at a specified tolerance level θ_{nt}. This probability is 1-β, where β is the probability of a false negative. Having determined this sample size we can then plot the probability of a positive result (calculated using eq. 1) for other levels of infection (θ) (Fig. 1). It is clear from Fig.1 that there is a considerable probability of rejecting seedlots with an infection level below the tolerance level. These are considered to be false positives and it is not until the infection level is around one tenth of the tolerance level that the probability of a false positive, α, is less than 0.25.

In the quarantine situation this does not present a problem as it is much more important to ensure that seedlots with an infection level above the threshold are rejected than it is to be concerned about the possibility of rejecting acceptable seedlots, i.e. α is no longer relevant. In practice this means that there is a reasonable margin of safety built into the testing programme.

In routine quality control, the problem is different, as it becomes more important not to reject acceptable seedlots. There is no way of reducing this margin for error as long as only one sample is tested. Therefore the benefits in terms of costs and simplicity of only testing a single sample must be balanced against the probability, α, of rejecting too many acceptable seed lots. Alternatively a lower probability, 1-β, for detecting the tolerance level could be accepted, reducing the number of seeds tested and shifting the line in Fig.1 to the right.

Case 2: Probability of detection <1

The previous case is based on the assumption that a single infected seed in the sample can always be detected. When this is not the case, the probability of obtaining a positive result (p_+) is multiplied by a factor p_s, the probability of detecting an infected seed in the sample, i.e $p_+ = p_s.p_c$ (Geng *et al.* 1983).

This problem can be dealt with simply by determining the maximum sample size in which an infected seed is almost certain to be detected and dividing the same number of seeds as determined above (Eq.3) into sub-samples of this size. A negative result from all such sub-samples is equivalent to obtaining a negative result for the original sample size. Thus, the criterion for rejecting a seedlot is: to reject it if one or more sub-samples gives a positive result. The probabilities of false positives, α, and false negatives, ß, are then precisely the same as if the sub-samples had been tested as one combined sample as in Case 1.

Alternatively, if the probability of detecting an infected seed, p_s, is known, it is possible to calculate the number of samples, k, of size n, which need to be tested to meet pre-defined tolerance levels using the following formula:

$$k = \frac{\ln(\beta)}{\ln[1 - p_s . p_c]} = \frac{\ln(\beta)}{\ln[1 - p_s(1-(1-\theta_{nt})^n)]} \quad \textbf{(4)}$$

However, it must be borne in mind (when the criterion for rejecting a seedlot is set as at least one sub-sample giving a positive result) that as the number of samples increases, the probability, α, of rejecting a seedlot with an infection level below the tolerance level increases (Table 2). When multiple samples are tested, it may be more appropriate to set the rejection/acceptance criterion to be something other than the rule of at least one positive sub-sample to reject a seedlot e.g. at least one negative sub-sample.

Another solution, especially if the probability of detecting an infected seed, p_s, is relatively close to one, would be to accept a lower probability for 1-ß, the probability of rejecting a seed lot with a non-tolerable level of infection.

Case 3: Quantitative estimation of the proportion of infected seeds

The simplest way of estimating the true proportion of infected seeds in a lot is to test each seed in a sample individually. The more seeds examined the more accurate the estimate becomes. When the proportion of infected seeds is small, however, this requires a very large number of individual seeds to be examined and becomes impractical. The solution is to test several samples of seeds and to estimate the proportion of infected seeds by the maximum likelihood method. The design problem is to determine the number and size of samples. The more samples tested, the more precise the estimate of infection (Table 3). Thus, the gain in precision must be balanced against the effort required to do the additional tests.

Determining the optimum sample size is more difficult. Little information is gained from an assay if either all the samples are positive or all are negative; an estimate of the infection level can only be made if some samples are positive and some are negative. It can be shown that most information is obtained from such an assay when the expected mean number of infected seeds per sample is between 1 and 2. Clearly the optimum sample size depends on the proportion of infected seeds, which is unknown. Therefore in order to set the sample size to be used, some *a priori* estimate of the infection level must be made. Although there may be little prior knowledge of infection levels, it may be possible to circumvent the problem by defining a range of interest.

There are two main strategies for dealing with this problem:

(a) test a number of samples of different sizes on one occasion (single stage design);
(b) perform a preliminary ranging test or tests prior to the main test (sequential or multi-stage design).

The choice of approach will depend on the aims of the assay, the precision required and the amount of prior knowledge. Thus, where the time taken for the final result is a major factor, as in routine tests on commercial seed, (a) may be the only option, but in research, where time may not be such an important factor, (b) may give the most efficient use of resources.

The problem of assay design in general terms has been considered by several workers (e.g. Finney, 1978; Strijbosch *et al.,* 1987) and more recently in the specific case of seed testing by Ridout (1993b).

Single stage design. Examples of optimum batch sizes, in the simplest but artificial case, where the infection level is known are given in Table 4 (from Ridout, 1993b). Generally the number of individual samples to be tested will be set in advance based on cost and the number of samples which can be reasonably be handled in the laboratory. The general conclusions are that when the range of the prior estimate is narrow, the samples should all be of the same size, but as the range of the prior estimate becomes wider then the samples should be of an increasing number of different sizes, but with some repetitions. For examples, see Table 5 (from Ridout 1993b).

Sequential design. As the optimum sample size can only be determined when the infection level is known in advance, it follows that if some preliminary or ranging tests can be done in advance then the final sample size to be used can be targeted more efficiently. The design problem is then to choose the size and number of samples to be used in the ranging tests. Intuitively, as the precision of the estimate increases with the number of samples tested, it would seem that if the total number of tests is constrained, most samples should be tested when most is known about the sample, i.e. if the test is done in two stages, more samples should tested at stage 2 than at stage 1.

Ridout (1993b) considered this problem with the constraints that a maximum of twelve samples should be tested in three stages and concluded that 1, 3 and 8 samples should be tested at stages 1, 2 and 3 respectively. Using this design and a prior range of the infection level between 0.1 and 25% the initial sample size was determined to be 26 seeds with subsequent sample sizes depending on the outcome of the tests at the previous stage (Table 6). The efficiency of this three stage approach was 70% compared to an efficiency of 40% for the optimal single stage design.

Interpretation. Previous formulae have presented the probability of a result occurring, given that we know the true infection level θ. The interpretation of the results is based on the reverse of this, i.e. what is the infection level, $\hat{\theta}$, which is most likely to have given the result obtained. Most simply this can be understood by examination of Table 7. Reading across the table gives the *probability* of a result occurring given certain infection level, θ; reading down the table gives the *likelihood* of the infection level being $\hat{\theta}$ given a certain result. Unlike probabilities which must sum to one, the absolute likelihood values have no intrinsic meaning and it is only their relative values which are important. It is therefore conventional to express them as ratios of the maximum value which can be obtained. Thus, we would divide the entries in Table 7 by the maximum of the appropriate column.

When tests have been done on a group of samples of the same size, and with p_s=1, the most likely proportion of infected seeds can be determined using the following formula:

$$\hat{\theta} = 1 - e^{\frac{\ln(1-\frac{r}{k})}{n}} \qquad \textbf{(5)}$$

where $\hat{\theta}$ is the infection level, n is the number of seeds in a sample and r/k is the proportion of positive results, r=number positive, k=number of samples. Effectively, the proportion of positive results, r/k, in this formula provides an estimate of the probability of a positive result, p_+.

When samples of several different sizes have been tested, although the same principle is involved, there is unfortunately no explicit formula as in equation 5. It is, of course, possible to obtain separate estimates for each sample size; however, a linear relationship between p_i, the probability of a positive result for the i^{th} sample, and θ can be obtained by rearrangement of equation (1):

$$\ln(-\ln(1-p_i)) = \ln(n_i) + \ln(-\ln(1-\theta)) \qquad \textbf{(6)}$$

which provides a basis for producing a combined estimate from all the data on a computer as in the procedure 'DILUTION' (Ridout & Welham, 1991) in the Genstat V procedure library. Equation 6 was presented graphically in a nomogram by Taylor and Phelps (1984) and Taylor *et al*. (1993). Traditionally, because of the computational requirements, estimates were presented in the form of 'MPN' tables as provided by Taylor and Phelps (1984) and Taylor *et al.* (1993). These are limited, however, because they only apply to defined testing schemes.

Confidence Limits. Calculation of confidence intervals for estimates has been the subject of some debate. Taylor (1970) considered that the 95% confidence intervals provided in MPN tables by Swaroop (1951) were unnecessarily wide and proposed the use of 80% limits calculated to Swaroop's tables. Taylor & Phelps (1984), in their tables, presented their estimates with "2 unit support limits" (Edwards, 1972), now generally known as likelihood ratio (LR) confidence intervals. Likelihood ratios follow an asymptotic χ^2 distribution (Cox & Snell,

1989) from which it follows that a likelihood ratio of 6.82 ($e^{\frac{\chi^2_{1,\alpha}}{2}}, \alpha = 0.05$ 7) is equivalent to a 95% confidence interval. In a theoretical investigation of the properties of confidence intervals estimated by different methods, Ridout (1993a) justified the use of confidence intervals based on the likelihood ratio for this problem. He noted that when all samples are either positive or negative only lower and upper limits, respectively, can be estimated and emphasised the importance of goodness-of-fit tests, as confidence intervals can become misleadingly small when the combination of results is improbable (i.e. the data has a poor fit to the model). The nomogram of Taylor & Phelps (1993) provides a visual measure of goodness-of-fit, but does not provide any confidence intervals. Unlikely combinations of results are often (but not always) omitted from MPN tables. The Genstat procedure 'DILUTION' (Ridout & Welham, 1991) provides both confidence intervals and a measure of goodness-of-fit.

Conclusions

It is apparent from the principles presented above that in designing any seed health assay, the aims of the assay should be clearly defined and that tolerance levels and what constitutes an acceptable confidence interval should be set in advance. Tolerance levels should be based on sound epidemiological data, and may vary depending on the aim of the assay.

The labour and cost requirements need to be carefully balanced in relation to the required standards of precision and reliability. For example, in one currently used scheme for soya bean bacterial blight testing, five samples of 1000 seeds are tested and seedlots are rejected only if all samples are positive. This implies a tolerance level of 0.44% for a probability level of 95%, the same tolerance level could be achieved by testing a single sample of 700 seeds and rejecting the seedlot if it is positive. This approach would increase the probability of rejecting acceptable seedlots compared to testing multiple samples, but presumably reduces the work (and cost) involved five-fold.

The sensitivity of the test method in terms of the probability of detecting a single infected seed in a sample of a certain size is of prime importance in determining whether more than one sample is needed, and it is vital that this is investigated when developing any test system. Clearly there is a need to define what is meant by an infected seed, in practice this could be defined as a seed carrying greater than a certain number of pathogen propagules. This, however, implies that there is some inoculum threshold below which disease will not occur and ignores the one-hit principle of infection which probably applies to most diseases (Roberts, 1985). It follows from the one-hit principle that a seed carrying a low number of disease propagules may have a very low probability of giving rise to disease, but this probability is not zero. Again this emphasises the need for a thorough understanding of the aetiology and epidemiology of the disease, particularly the relationship between inoculum, disease (transmission), and detection thresholds.

Throughout this paper, it has been assumed that the seeds tested are a random sample from the seedlot and that the sampling protocols are adequate to ensure this. These assumptions are of prime importance and it is unlikely that any assay can compensate for poor or biased sampling of the seed bulk.

Computer Programs

A 'stand alone' computer program to estimate the proportion of infected seeds from the results of an assay, together with confidence intervals and a goodness-of-fit test, will soon be available from the authors.

Acknowledgements

This work was funded by the Ministry of Agriculture, Fisheries and Food (England and Wales).

References

Cochran, W.G. (1950) Estimation of bacterial densities by means of the "Most Probable Number". *Biometrics,* **6,** 105-116.

Cox, D.R. & Snell, E.J. (1989) *Analysis of Binary Data.* London: Chapman and Hall.

Geng, S., Campbell, R.N., Carter, M. & Hills, M. (1983) Quality control programs for seedborne pathogens. *Plant Disease* **67,** 236-242.

Marrou, J. & Messiaen, C.M. (1967) The *Chenopodium quinoa* test: a critical method for detecting seed transmission of lettuce mosaic virus. *Proceedings of the International Seed Testing Association*, **32**, 49-57.

Ridout, M.S. & Welham, S (1991) DILUTION: Analysis of dilution series data. In: eds. Payne, R.W. and Arnold, G.M. *Genstat V Procedure Library, Release 2[2]*. Oxford: Numerical Algorithms Group. pp.88-90.

Ridout, M.S. (1993a) A comparison of confidence interval methods for dilution series experiments. *Biometrics* (in press)

Ridout, M.S. (1993b) Three-stage designs for seed testing experiments. *Applied Statistics* (in press)

Roberts, S.J. (1985) A note on *Pseudomonas syringae* pv. *berberidis* infections of *Berberis*: comparative infection studies in attached and detached leaves. *Journal of Applied Bacteriology*, **59,** 369-374.

Russell, T.S. (1988) Some aspects of sampling and statistics in seed health testing and the establishment of threshold levels. *Phytopathology*, **78,** 880-881.

Swaroop, S. (1951) The range of variation in the most probable number of organisms estimated by the dilution method. *Indian Journal of medical Research*, **39**, 107-134.

Strijbosch, L.W.G., Buurman, W.A., Does, R.J.M.M., Zinken, P.H. & Grönewegen, G. (1987) Limiting dilution assays. Experimental design and statistical analysis. *Journal of Immunological Methods*, **97**, 133-140.

Taswell, C. (1981) Limiting dilution assays for the determination of immunocompetent cell frequencies. *Journal of Immunology*, **126**, 1614-1619.

Taylor, J.D. (1970) The quantitative estimation of the infection of bean seed with *Pseudomonas phaseolicola* (Burkh.) Dowson. *Annals of Applied Biology* **66,** 29-36.

Taylor, J.D & Phelps, K. (1984) Basic method for the detection of seed-borne bacteria 3. Estimation of percentage seed infection. *Report of the 1st International Workshop on Seed Bacteriology 4-9 October 1982. Angers, France.* Zurich: ISTA Secretariat. pp.12-14.

Taylor, J.D., Phelps, K. & Roberts, S.J. (1993) Most probable number (MPN) method: origin and application. *this conference.*

Table 1. Samples sizes, *n*, for probability of 0.95 or 0.99 of a positive test result for an infection level, θ.

θ	*n* (β≈0.05)	*n* (β≈0.01)
0.01%	29960	46050
0.02%	14980	23020
0.05%	5990	9210
0.1%	2990	4600
0.2%	1500	2300
0.5%	600	920
1%	300	460
20%	150	230
50%	4	10

Table 2. Some examples of the number of samples, *k*, required for a probability of 0.95 of detection of an infection level θ_{nt}=0.1%, given a test sensitivity of p_s and a sample size of *n*, and probability of accepting a seed lot, 1-α, with a tolerable infection level of θ_t=0.05%

p_s	*n*	Rounded *k* (1-ß=0.95, θ_{nt}=0.1%)	1-α (θ_t=0.05%)
1	3000	1	0.22
1	1000	3	0.22
0.8	3000	2	0.08
0.96	1000	3	0.24
1	30	100	0.22
1	100	30	0.22
0.87	200	17	0.23
0.65	500	10	0.19

Table 3. Effect of number of samples, *k*, on 95% confidence limits for a sample size, *n*, of 1000 seeds.

Number of samples, *k*	Number positive, *r*	Estimated infection level, (%)	95 % confidence limits	
			lower	upper
3	2	0.11	0.02	0.38
6	4	0.11	0.03	0.27
12	8	0.11	0.05	0.21
24	16	0.11	0.06	0.18
48	32	0.11	0.07	0.16

Table 4. Optimum sample size, *n*, for determining infection level, θ.

θ	*n*
0.001	1593
0.005	318
0.01	159
0.05	31
0.10	15
0.25	6

Table 5. Sample sizes for optimal single stage seed testing design for different prior ranges and testing six samples (from Ridout, 1993b).

Prior range	Sample sizes
0.4-6.7%	64,64,64,64,64,64
0.1-25%	16,16,16,89,303,303
0.006%-99%	1,1,14,102,755,5202

Table 6. Sample sizes for optimal three stage seed testing design for a prior range of infection levels from 0.1 to 25% and testing a total of twelve samples on three occasions (from Ridout, 1993b).

Stage 1		Stage 2		Stage 3
Size	Outcome	Size	Outcome	Size
26	0/1	240	0/3	902
			1/3	639
			2/3	335
			3/3	65
	1/1	13	0/3	95
			1/3	35
			2/3	16
			3/3	9

Table 7. Probabilities of obtaining *r* positive samples out of 5 tested for a sample size, *n*, of 1 and an infection level of θ

θ	Number of positive samples, *r*, out of 5 tested					
	0	1	2	3	4	5
5%	0.77	0.20	0.02	0.00	0.00	0.00
10%	0.59	0.33	0.07	0.01	0.00	0.00
20%	0.33	0.41	0.20	0.05	0.01	0.00
30%	0.17	0.36	0.31	0.13	0.03	0.00
40%	0.08	0.26	0.35	0.23	0.08	0.01
50%	0.03	0.16	0.31	0.31	0.16	0.03
60%	0.01	0.08	0.23	0.35	0.26	0.08
70%	0.00	0.03	0.13	0.31	0.36	0.17
80%	0.00	0.01	0.05	0.20	0.41	0.33
90%	0.00	0.00	0.01	0.07	0.33	0.59
95%	0.00	0.00	0.00	0.02	0.20	0.77

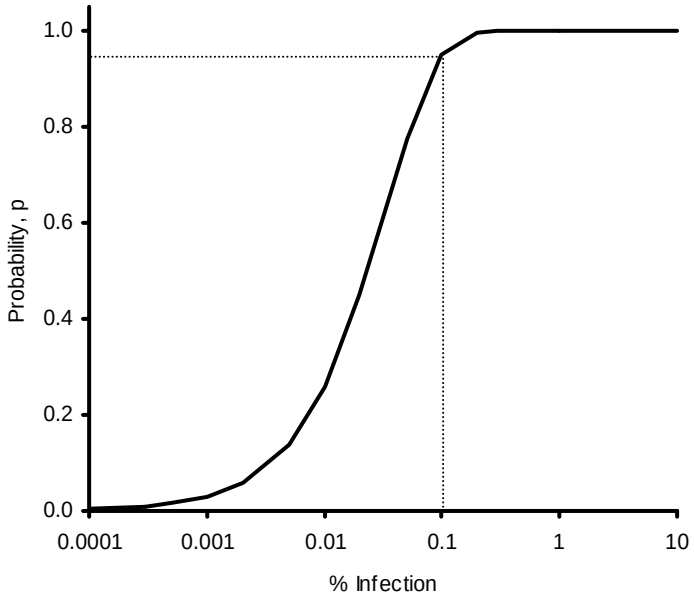

Figure 1. Probability, p_c of at least one infected seed being contained in a sample of 3,000 seeds for different infection levels, θ, calculated using the formula: $p_c = 1 - (1 - \theta)^n$.